arXiv:1511.08501v1 [stat.ME] 26 Nov 2015

Variable Selection in Causal Inference using a Simultaneous Penalization Method

BY ASHKAN ERTEFAIE[1], MASOUD ASGHARIAN[2] AND DAVID A. STEPHENS[2]
[1]*Department of Statistics, University of Pennsylvania, Philadelphia, PA, USA*
[2]*Department of Mathematics and Statistics, McGill University, Montreal, Canada*
ertefaie@wharton.upenn.edu

SUMMARY

In the causal adjustment setting, variable selection techniques based on one of either the outcome or treatment allocation model can result in the omission of confounders, which leads to bias, or the inclusion of spurious variables, which leads to variance inflation, in the propensity score. We propose a variable selection method based on a penalized objective function which considers the outcome and treatment assignment models simultaneously. The proposed method facilitates confounder selection in high-dimensional settings. We show that under regularity conditions our method attains the oracle property. The selected variables are used to form a doubly robust regression estimator of the treatment effect. We show that under some conditions our method attains the oracle property. Simulation results are presented and economic growth data are analyzed. Specifically, we study the effect of life expectancy as a measure of population health on the average growth rate of gross domestic product per capita.

Some key words: Causal inference, Average treatment effect, Propensity score, Variable selection, Penalized estimator, Oracle estimator.

1. INTRODUCTION

In the analysis of observational data, when attempting to establish the magnitude of the causal effect of treatment (or exposure) in the presence of confounding, the practitioner is faced with certain modeling decisions that facilitate estimation. Should one take the parametric approach, at least one of two statistical models must be proposed; (i) the *conditional mean model* that models the expected outcome as a function of predictors, and (ii) the *treatment allocation model* that describes the mechanism via which treatment is allocated to (or, at least, received by) individuals in the study, again as a function of the predictors (Rosenbaum & Rubin, 1983; Robins & Brumback, 2000).

Predictors that appear in both mechanisms (i) and (ii) are termed *confounders*, and their omission from model (ii) is typically regarded as a serious error, as it leads to inconsistent estimators of the treatment effect. Thus practitioners usually adopt a conservative approach, and attempt to ensure that they do not omit confounders by fitting a richly parameterized treatment allocation model. The conservative approach, however, can lead to predictors of treatment allocation only – and not outcome – being included in the treatment allocation model. The inclusion of such "spurious" *instrumental* variables in model (ii) is usually regarded as harmless. However, the typical forfeit for this conservatism is inflation of variance of the effect estimator (Greenland, 2008; Schisterman et al., 2009). This problem also applies to the conditional mean model, but is in practice less problematic, as practitioners seem to be more concerned with bias removal, and

therefore more likely to introduce the spurious variables in model (ii). Little formal guidance as to how the practitioner should act in this setting has been provided.

As has been conjectured and studied in simulation by Brookhart et al. (2006a), it is plausible that judicious variable selection may lead to appreciable efficiency gains, and several approaches with this aim have been proposed. However, confounder selection methods based on either just the treatment assignment model or just the outcome model may fail to account for non-ignorable confounders which barely predict the treatment or the outcome, respectively (Crainiceanu et al., 2008; Belloni et al., 2014): in this manuscript, we use the term *weak confounder* for these variables. Vansteelandt et al. (2010) shows that confounder selection procedures based on AIC and BIC can be sub-optimal and introduce a method based on the focused information criterion (FIC) which targets the treatment effect by minimizing a prediction mean square error (see also the cross-validation method of Brookhart & van der Laan (2006b)). Van der Laan et al. (2007) introduces a *Super Learner* estimator which is computed by selecting a candidate from a set of estimators obtained from different models using a cross-validation risk (Van der Laan et al., 2004; Sinisi et al., 2007).

Bayesian adjustment for confounding (BAC) is a parametric variable selection approach introduced by Wang et al. (2012); BAC specifies a prior distribution for a set of possible models which includes a dependence parameter, $w \in [1, \infty]$, representing the odds of including a variable in the outcome model given that the same variable is in the propensity score model. If we *know* a priori that a predictor of treatment is in fact a confounder, then w can be set to ∞ (Crainiceanu et al., 2008; Zigler et al., 2013). Wilson & Reich (2014) proposes a decision-theoretic approach to confounder selection that can handle high-dimensional cases; a Bayesian regression model is fit and using the posterior credible region of the regression parameters, a set of candidate models is formed. A sparse model is then found by penalizing models that do not include confounders. This method is conservative in the sense that it may include instrumental variables that may inflate the variance of the treatment effect. Also, tuning the penalty function can be challenging.

Asymptotically, it is known that penalizing the conditional outcome model, given treatment and covariates, results in a valid variable selection strategy for causal effect estimation; however, for small to moderate sample sizes, it may result in the omission of weak confounders. The objective of this manuscript is to improve the small sample performance of the outcome penalization strategy while maintaining its asymptotic performance (Table 6). We present a covariate selection procedure which facilitates the estimation of the treatment effect in the high-dimensional cases. Specifically, we propose a penalized objective function which considers both covariate-treatment and covariate-outcome associations and has the ability to select even weak confounders. This objective function is used to identify the set of non-ignorable confounders and predictors of outcome; the resulting parameter estimates do not have any causal interpretation. We derive the asymptotic properties of procedure and show that under some mild conditions the estimators have oracle properties (specifically, are consistent and asymptotically normally distributed). We utilize the selected covariates to estimate the causal effect of interest using a doubly robust estimator.

2. PRELIMINARIES & NOTATION

In standard notation, let $Y(d)$ denote the (potential) outcome arising from treatment d, and let D denote the treatment received. We consider for illustration the case of binary treatment. The observed outcome, Y, is defined as $DY(1) + (1 - D)Y(0)$. We restrict attention here to the situation where each predictor can be classified into one of three types, and to single time-point studies. We consider

Fig. 1: Covariate types: Type-I: X_1, Type-II: X_2 and Type-III: X_3.

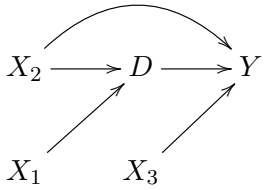

(I) *treatment predictors* (X_1), which are related to treatment and not to outcome.
(II) *confounders* (X_2), which are related to both outcome and treatment.
(III) *outcome predictors* (X_3), which are related to outcome and not to treatment;

see the directed acyclic graph (DAG) in Figure 1. In addition, as is usual, we will make the assumption of *no unmeasured confounders*, that is, that treatment received D and potential outcome to treatment d, $Y(d)$, are independent, given the measured predictors X. In any practical situation, to facilitate causal inference, the analyst must make an assessment as to the structural nature of the relationships between the variables encoded by the DAG in Figure 1.

2·1. *The Propensity Score for binary treatments*

The *propensity score*, $\pi(.)$, for binary treatment D is defined as $\pi(x) = \Pr(D = 1|x)$, where x is a p-dimensional vector of (all) covariates. In its random variable form, Rosenbaum & Rubin (1983) show that $\pi(X)$ is the coarsest function of covariates that exhibits the balancing property, that is, $D \perp X|\pi(X)$. As a consequence, the causal effect $\mu = \mathbb{E}[Y(1) - Y(0)]$ can be computed by iterated expectation

$$\mu = \mathbb{E}_X[\mathbb{E}\{Y(1)|X\} - \mathbb{E}\{Y(0)|X\}] = \mathbb{E}_\pi[\mathbb{E}\{Y(1)|\pi\} - \mathbb{E}\{Y(0)|\pi\}], \tag{1}$$

where $\mathbb{E}_\pi$ denotes the expectation with respect to the distribution of $\pi(X)$. For more details see Rubin (2008) and Rosenbaum (2010).

Remark 1: In the standard formulation of the propensity score, no distinction is made between our three types of covariates. Note that, however, for consistent estimation of μ, *it is not necessary to balance on covariates that are not confounders*. Covariates X_1 that predict D but not Y may be unbalanced in treated and untreated groups, but will not affect the estimation of the effect of D on Y, as D will be conditioned upon, thereby blocking any effect of X_1 (De Luna et al., 2011). Covariates X_3 are unrelated to D, so will by assumption be in balance in treated and untreated groups in the population. Therefore, the propensity score need only be constructed from confounding variables X_2. Then, in the presence of outcome predictors X_3 of Y, the sequel to equation (1) takes the form

$$\begin{aligned}\mu = \mathbb{E}[Y(1) - Y(0)] &= \mathbb{E}_{X_2,X_3}[\mathbb{E}\{Y(1)|X_2, X_3\} - \mathbb{E}\{Y(0)|X_2, X_3\}] \\ &= \mathbb{E}_{\pi_2,X_3}[\mathbb{E}\{Y(1)|\pi_2, X_3\} - \mathbb{E}\{Y(0)|\pi_2, X_3\}].\end{aligned} \tag{2}$$

Remark 2: Inclusion of covariates that are just related to the outcome in the propensity score model increases the covariance between the fitted π and Y, and decreases the variance of the estimated causal effect, in line with the simulation of Brookhart et al. (2006a).

2·2. *Penalized Estimation*

In a given parametric model, if α is a r-dimensional regression coefficient, $p_\lambda(.)$ is a penalty function and $l_m(\alpha)$ is the negative log-likelihood, the maximum penalized likelihood (MPL) estimator $\widehat{\alpha}_{ml}$ is defined as

$$\widehat{\alpha}_{ml} = \arg\min_{\alpha} \left[l_m(\alpha) + n \sum_{j=1}^{r} p_\lambda(|\alpha_j|) \right].$$

MPL estimators are shrinkage estimators, and as such, they typically have more finite sample bias, though less variation than unpenalized analogues. Commonly used penalty functions include LASSO (Tibshirani, 1996), SCAD (Fan & Li, 2001), Elastic Net (EN) (Zou & Hastie, 2005) and HARD (Antoniadis, 1997).

The remainder of this paper is organized as follows. Section 3 presents our two step variable selection and estimation procedure; we establish its theoretical properties. The performance of the proposed method is studied via simulation in Section 4. We analyze a real data set in Section 5, and Section 6 contains concluding remarks. All proofs are given in the Appendix.

3. PENALIZATION AND TREATMENT EFFECT ESTIMATION

In this section, we present our proposed method for estimating the treatment effect in high-dimensional cases. We separate the covariate selection and treatment effect estimation procedure. First, we form a penalized objective function which is used to identify the important covariates, and establish the theoretical properties of the resulting estimators (i.e., minimizers of the penalized objective function). Note that because of the special characteristics of this function the estimators do not have any causal interpretation and are used just to *prioritize* variables. Second, treatment effect estimation is performed using a doubly robust estimator with the selected covariates. We use a simple model structure to illustrate the methodology, and assume that a random sample of size n of observations of outcome, exposure and covariates is available.

In order to present the method, we initially assume that columns of $\mathbf{X}$ are orthogonal; this simplifying assumption is relaxed in Section 3·7.

3·1. *Penalized objective function*

Consider the following linear outcome model under a binary exposure d

$$Y = \theta d + \mathbf{x}\alpha_y + \epsilon,$$

where $\mathbf{x}$ is a $1 \times r$ standardized covariate vector, and ϵ is standard normal residual error. Assuming a logit model for the propensity score, we have

$$\pi(\mathbf{x}, \alpha_d) = p(D = 1|\mathbf{x}, \alpha_d) = \frac{\exp\{\mathbf{x}\alpha_d\}}{1 + \exp\{\mathbf{x}\alpha_d\}}.$$

Let $\mathbf{X}$ denote the corresponding $n \times r$ design matrix. In the formulation, α_y and α_d denote $r \times 1$ vectors of parameters and θ is the treatment effect.

First, we form an objective function, $M(\alpha)$, in which the coefficients α_y and α_d are replaced by a single common vector of coefficients α that is proportional to a weighted sum of $|\alpha_y|$ and $|\alpha_d|$, where $|.|$ denotes the componentwise absolute value. This will guarantee that the objective function satisfies the following condition:

$Argmin$ Condition: An element, $\hat{\alpha}_j$, of the minimizer $\hat{\alpha}$ of $M(\alpha)$ converges to zero as $n \to \infty$ if the corresponding covariate is not associated with either treatment or outcome.

In standard linear regression where we regress Y on a vector of covariates $\mathbf{X}$ (with no treatment variable) that are presumed orthogonal, an example of an objective function that estimates $|\alpha_y|$ is

$$M(\alpha) = \frac{1}{2n}\left(\left|\sum_{i=1}^{n}\mathbf{x}_i^\top y_i\right| - n\alpha\right)^\top \left(\left|\sum_{i=1}^{n}\mathbf{x}_i^\top y_i\right| - n\alpha\right)$$

yielding the estimating equation and estimate

$$\frac{\partial M(\alpha)}{\partial \alpha} = \left|\sum_{i=1}^{n}\mathbf{x}_i^\top y_i\right| - n\alpha = 0 \qquad \Longrightarrow \qquad \widehat{\alpha} = \frac{1}{n}\left|\sum_{i=1}^{n}\mathbf{x}_i^\top y_i\right|.$$

We have used the fact $n^{-1}\mathbf{x}^\top\mathbf{x}$ is an identity matrix due to standardization. It is clear that as $n \longrightarrow \infty$, $\widehat{\alpha} \xrightarrow{p} |\mathbb{E}[x_i^\top Y_i]|$ Similarly, in a binary exposure setting, $|\alpha_d|$ can be estimated using the objective function $[-|\sum_{i=1}^{n}\mathbf{x}_i d_i|\,\alpha + \sum_{i=1}^{n}\log(1+\exp\{\mathbf{x}_i\alpha\})]$. Now, we combine these two objective functions and show that the resulting objective function have some interesting features. First, we obtain the least squares (or ridge) estimate $\widetilde{\theta}$ of θ by regressing Y on D and the vector of covariates $\mathbf{X}$. Because of the inclusion of spurious and instrumental variables in the model, $\widetilde{\theta}$ is not efficient but is consistent for θ under the no unmeasured confounders assumption. Define $\widetilde{Y} = Y - \widetilde{\theta}D$. Let

$$M(\alpha) = \frac{1}{2}\Big[\Big|\sum_{i=1}^{n}\mathbf{x}_i\widetilde{y}_i\Big| - \alpha^\top\sum_{i=1}^{n}(\mathbf{x}_i^\top\mathbf{x}_i)\Big]\left(\sum_{i=1}^{n}\mathbf{x}_i^\top\mathbf{x}_i\right)^{-1}\Big[\Big|\sum_{i=1}^{n}\mathbf{x}_i\widetilde{y}_i\Big| - \alpha^\top\sum_{i=1}^{n}(\mathbf{x}_i^\top\mathbf{x}_i)\Big]^\top + \frac{1}{\tau}\Big[-\Big|\sum_{i=1}^{n}\mathbf{x}_i d_i\Big|\alpha + \sum_{i=1}^{n}\log(1+\exp\{\mathbf{x}_i\alpha\})\Big], \quad (3)$$

where τ is a positive constant. Under orthogonality and standardization of $\mathbf{X}$, we have

$$M(\alpha) = \frac{1}{2n}\Big[\Big|\sum_{i=1}^{n}\mathbf{x}_i\widetilde{y}_i\Big| - n\alpha^\top\Big]\Big[\Big|\sum_{i=1}^{n}\mathbf{x}_i\widetilde{y}_i\Big| - n\alpha^\top\Big]^\top + \frac{1}{\tau}\Big[-\Big|\sum_{i=1}^{n}\mathbf{x}_i d_i\Big|\alpha + \sum_{i=1}^{n}\log(1+\exp\{\mathbf{x}_i\alpha\})\Big]$$

Note that for each j, the parameter α_j corresponding to x_j is the same in both models. Now, we show that in this function the absolute values play a critical rule in satisfying the $Argmin$ condition. Let $\hat{\alpha} = \arg\min_\alpha M(\alpha)$ and $\widetilde{\alpha}_y$ be the least squares estimate of the parameters in the outcome model. By convexity of $M(\alpha)$, $\hat{\alpha}$ must be a solution to $\partial M(\alpha)/\partial\alpha = 0$, which implies that

$$\left[n\widehat{\alpha} + \frac{1}{\tau}\sum_{i=1}^{n}\mathbf{x}_i^\top\frac{\exp\{\mathbf{x}_i\widehat{\alpha}\}}{1+\exp\{\mathbf{x}_i\hat{\alpha}\}}\right] = \left|\sum_{i=1}^{n}\mathbf{x}_i^\top\widetilde{y}_i\right| + \frac{1}{\tau}\left|\sum_{i=1}^{n}\mathbf{x}_i^\top d_i\right| = n\,|\widetilde{\alpha}_y| + \frac{1}{\tau}\left|\sum_{i=1}^{n}\mathbf{x}_i^\top d_i\right|. \quad (4)$$

This equation shows that $\hat{\alpha} = 0$ (i.e. the $Argmin$ condition is satisfied) if $cov(X,Y) = cov(X,D) = \mathbf{0}$. Note that $cov(X,D) = \mathbf{0}$ implies $cov(X,\widetilde{Y}) = cov(X,Y)$. Moreover, using

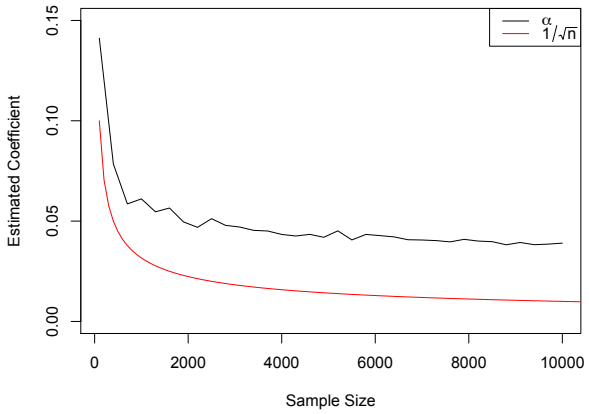

Fig. 2: Performance of the modified objective function based estimator for different sample sizes n. Red and black lines are $1/\sqrt{n}$ and the estimated coefficient $\widehat{\alpha}$ using the modified objective function.

the first three terms of a Taylor expansion, we have

$$\widehat{\alpha} = \frac{2\tau}{2\tau+1}\left|\widetilde{\alpha}_y\right| + \frac{1}{2\tau+1}\left|\widetilde{\alpha}_d\right|.$$

Thus, $\hat{\alpha}$ is a weighted sum of $|\widetilde{\alpha}_y|$ and $|\widetilde{\alpha}_d|$ where $\widetilde{\alpha}_d$ is the maximum likelihood estimate of the parameters in the treatment model. The constant τ controls the contribution of components $|\widetilde{\alpha}_y|$ and $|\widetilde{\alpha}_d|$ to the estimator. For example, for $\tau = 2$, $\widehat{\alpha} = \frac{4}{5}\left|\widetilde{\alpha}_y\right| + \frac{1}{5}\left|\widetilde{\alpha}_d\right|$; while for $\tau = 0.1$, $\widehat{\alpha} = \frac{0.2}{1.2}\left|\widetilde{\alpha}_y\right| + \frac{1}{1.2}\left|\widetilde{\alpha}_d\right|$. Thus, τ gives a flexibility to our objective function such that as it decreases to zero the proposed estimate $\widehat{\alpha}$ converges to $|\widetilde{\alpha}_d|$. See Section 3·2.

Figure 2 visually presents how $\widehat{\alpha}_j$ behaves, for a fixed $\tau = 0.5$, when $\widehat{\alpha}_{jy}$ converges to zero as sample size increases, i.e., the jth variable becomes insignificant. Specifically, this figure presents a case where there is just one covariate (i.e., $j = 1$) and the coefficient of this covariate in outcome and treatment models are $\alpha_{jy} = 1/\sqrt{n}$ and $\alpha_{jd} = 0.3$, respectively, where n is the sample size. As expected, $\hat{\alpha}$ corresponding to this covariate does not converge to zero as sample size increases. The same behaviour would be observed if $\alpha_{jd} \to 0$ as $n \to \infty$ and α_{jy} is a non-zero constant.

Therefore, penalizing objective function (3) results in selecting covariates that are either related to outcome (Type-III) or treatment (Type-I). However, this is against our goal of keeping variables that are either predictors of the outcome or non-ignorable confounders and excluding instrumental variables (Type-I). To deal with this problem, we present a weighted lasso penalty function that is tailored specifically for causal inference variable selection:

$$\frac{\lambda}{(\widetilde{\alpha}_y)^2(1+|\widetilde{\alpha}_d|)^2}\sum_{j=1}^{r}|\alpha_j|,$$

where $\widetilde{\alpha}_y$ and $\widetilde{\alpha}_d$ are the least squares (or ridge) and maximum likelihood estimates of the parameters in the outcome and treatment models, respectively, and λ is a tuning parameter. Thus,

the proposed modified penalized objective function is given by

$$M_p(\alpha) = M(\alpha) + \lambda \sum_{j=1}^{r} \frac{|\alpha_j|}{(\widetilde{\alpha}_{jy})^2(1+|\widetilde{\alpha}_{jd}|)^2} \tag{5}$$

We refer to $\widehat{\alpha} = \arg\min_\alpha M_p(\alpha)$ as penalized modified objective function estimators (PMOE). The magnitude of the penalty on each parameter is proportional to its contribution to the outcome and treatment model. Note that as $\hat{\alpha}_{jy} \to 0$, our penalty function puts more penalty on the jth parameter while considering the covariate-treatment association. For example, when a covariate barely predicts the outcome and treatment, our proposed penalty function imposes a stronger penalty on the parameter compared to a case where a covariate barely predicts the outcome and is strongly related to treatment. This is an important feature of the proposed penalty function which allows selecting such weak confounders for small sample sizes. The proposed penalized estimator asymptotically selects the same covariates as the adaptive lasso on the outcome model (Zou, 2006). Thus, the proposed method improves the small sample performance of the outcome model penalization strategy while maintaining the asymptotic properties of this strategy.

3·2. *The role of τ in PMOE*

The constant τ reflects investigators belief about the importance of including variables that are weakly(strongly) related to outcome and strongly(weakly) related to treatment, i.e., for smaller values of τ, variables that are weakly related to the outcome but strongly to the treatment have more contribution to $\widehat{\alpha}$ and have more chance to be selected. Also, when $\tau = \infty$, the treatment mechanism does not have any contribution in the objective function (3). Thus, our procedure performs similarly to the outcome model based variable selection that may exclude non-ignorable confounders that are weakly related to the outcome. Our simulation studies shed more light on the role of τ in our procedure.

3·3. *Main Theorem*

Suppose $\alpha_0 = (\alpha_{01}, \alpha_{02})$ is the true parameter value of the r-dimensional vector of parameters where $\alpha_{02} = \{\alpha_j, j = s+1, ..., r\} \equiv \mathbf{0}$ contains those elements of α that are in fact zero, so that the corresponding predictors are not confounders; s denotes the true number of predictors present in the model (*exact sparsity* assumption). Let $\widehat{\alpha} = (\widehat{\alpha}_1, \widehat{\alpha}_2)$ be the vector of estimators corresponding to (5). The next theorem proves the sparsity and asymptotic normality of the proposed penalized estimators under the following two conditions:

(a) Let

$$\sum_{i=1}^{n} \epsilon_{i\widetilde{y}} = |\sum_{i=1}^{n} \mathbf{x}_i \widetilde{y}_i| - n\alpha_0^\top \qquad \text{and} \qquad \sum_{i=1}^{n} \epsilon_{id} = |\sum_{i=1}^{n} d_i \mathbf{x}_i| - \sum_{i=1}^{n} \mathbf{x}_i \frac{\exp\{\mathbf{x}_i\alpha_0\}}{1+\exp\{\mathbf{x}_i\alpha_0\}}$$

and $e_i = \epsilon_{i\widetilde{y}} + \frac{1}{\tau}\epsilon_{id}$. We assume that $\frac{1}{\sqrt{n}}\sum_{i=1}^{n}[\epsilon_{i\widetilde{y}} + \frac{1}{\tau}\epsilon_{id}]$ converges in distribution a multivariate normal $N_r(0, \Sigma(\alpha_0))$.

(b) Let

$$\frac{1}{n}\left[n\mathbb{1} + \frac{1}{\tau}\sum_{i=1}^{n} \frac{\exp\{\mathbf{x}_i\alpha_0\}}{[1+\exp\{\mathbf{x}_i\alpha_0\}]^2}\mathbf{x}_i^\top \mathbf{x}_i\right] \xrightarrow{p} \Omega(\alpha_0),$$

where $\Omega(\alpha_0)$ is a $r \times r$ positive definite matrix and $\mathbb{1}$ is the identity matrix.

THEOREM 1. ***(Oracle properties)*** *Suppose conditions (a) & (b) are fulfilled, further* $\lambda_n/\sqrt{n} \to 0$ *and* $\lambda_n\sqrt{n} \to \infty$*. Then*

(a) $\Pr(\widehat{\alpha}_2 = \mathbf{0}) \to 1$ *as* $n \to \infty$

(b) $\sqrt{n}(\widehat{\alpha}_{01} - \alpha_{01}) \xrightarrow{d} N(0, \Omega_{11}^{-1}\Sigma_{11}\Omega_{11}^{-1})$,

where $\alpha_{01} = \alpha_{01}$ *is the true vector of non-zero coefficients. Also,* Ω_{11} *and* Σ_{11} *are corresponding elements of* Ω *and* Σ*, respectively.*

3·4. *Choosing the Tuning Parameter*

We select the tuning parameter using the *Generalized Cross Validation* (GCV) method suggested by Tibshirani (1996) and Fan & Li (2001). Let $\mathbf{X}_\lambda$ be the selected covariates corresponding to a specific value of λ, we first regress $\widetilde{Y}$ on $\mathbf{X}_\lambda$ and calculate the residual sum of square (RSS) of this model. Then

$$\text{GCV}(\lambda) = \frac{\text{RSS}(\lambda)/n}{[1 - d(\lambda)/n]^2},$$

where $d(\lambda) = \text{trace}[\mathbf{X}_\lambda(\mathbf{X}_\lambda^\top\mathbf{X}_\lambda + n\Sigma_\lambda(\widehat{\alpha}))^{-1})\mathbf{X}_\lambda^\top]$ is the effective number of parameters and $\Sigma_\lambda(\alpha) = \text{diag}[p'_\lambda(|\alpha_1|)/|\alpha_1|, ..., p'_\lambda(|\alpha_r|)/|\alpha_r|]$. The selected tuning parameter $\widehat{\lambda}$ is defined by $\widehat{\lambda} = \arg\min_\lambda \text{GCV}(\lambda)$.

3·5. *Estimation of the causal effect*

For treatment effect estimation, we fit the following model using the set of covariates selected in the previous step; note that a user may want to use other causal adjustment models such as inverse probability weighting or propensity score matching, and the selection approach can be used for these procedures also.

Our model is a slight modification of the conventional propensity score regression approach of Robins et al. (1992), and specifies

$$\mathbb{E}[Y_i|S_i = s_i, \mathbf{X}_i = \mathbf{x}_i] = \theta s_i + g(\mathbf{x};\gamma), \tag{6}$$

where $S_i = D_i - \mathbb{E}[D_i|\mathbf{x}_i] = D_i - \pi(\mathbf{x}_i)$, $g(\mathbf{x};\gamma)$ is a function of covariates and π is the propensity score. The quantity S_i is used in place of D_i; if D_i is used the fitted model may result in a biased estimator for θ since $g(\mathbf{x};\gamma)$ may be incorrectly specified. By defining S_i in this way, we restore $cor(S_i, \mathbf{X}_{ij}) = 0$ for $j = 1, 2, .., p$ where p is the number of selected variables (if $\pi(\mathbf{x}_i) = \mathbb{E}[D_i|\mathbf{x}_i]$ is correctly specified), as $\pi(\mathbf{x}_i)$ is the (fitted) expected value of D_i, and hence $\mathbf{x}_j^\top(D - \pi(\mathbf{x})) = 0$, where $\mathbf{x}_j^\top = (x_{1j}, \ldots, x_{nj})$. Therefore, misspecification of $g(.)$ will not result in an inconsistent estimator of θ.

In general, this model results in a *doubly robust* estimator (see Davidian et al. (2005), Schafer & Kang (2005) and Bang & Robins (2005)); it yields a consistent estimator of θ if *either* the propensity score model or conditional mean model (6) is correctly specified, and is the most efficient estimator (Tsiatis, 2006) when both are correctly specified. For additional details on the related asymptotic and finite sample behavior, see Kang & Schafer (2007), Neugebauer & van der Laan (2005), van der Laan & Robins (2003) and Robins (2000).

Remark 3: The importance of the doubly robust estimator is that, provided the postulated outcome and treatment model identify the true non-zero coefficients in each model, the proposed estimation procedure will consistently estimate the treatment effect given that at least one of the propensity score or outcome models are correctly specified. Assuming linear working models, a

sufficient but *not* necessary condition for selecting non-ignorable confounders is the linearity of the true models in their parameters.

The model chosen for estimation of the treatment effect is data dependent. Owing to the inherited uncertainty in the selected model, making statistical inference about the treatment effect becomes "post-selection inference". Hence, inference about the treatment effect obtained in the estimation step needs to be done with caution. The weak consistency of the post-selection estimator results from the following theorem.

THEOREM 2. *Let $\zeta(\hat{\theta}_{M_n}, M_n)$ be a smooth function of $\hat{\theta}_{M_n}$ where M_n denotes a model formed from a set of selected variables using our method. Then $\zeta(\hat{\theta}_{M_n}, M_n) \xrightarrow{p} \zeta(\theta_{M_0}, M_0)$ as $n \to \infty$ where M_0 denotes a model formed from the set of non-zero coefficients.*

3·6. *The Procedure Summary*

The penalized treatment effect estimation process explained in Sections 3.1 to 3·5 can be summarized as follows:

1. Estimate the vector of parameters $\hat{\alpha} = \arg\min_\alpha M_p(\alpha)$ where $M_p(\alpha)$ is defined in (5).
2. Using the covariates with $\alpha \neq 0$, estimate the propensity score $\hat{\pi}(\mathbf{X})$.
3. Define a random variable $S_i = D_i - \hat{\pi}(\mathbf{X}_i)$ and fit the outcome model $\mathbb{E}[Y_i|d, \mathbf{x}] = \theta s_i + g(\mathbf{x}_i; \gamma)$. The vector of parameters (θ, γ) is estimated using ordinary least squares. For simplicity, we assume the linear working model for $g(\mathbf{x}_i; \gamma) = \gamma^\top \mathbf{x}_i$. The design matrix $\mathbf{X}$ includes a subset of variables with $\alpha \neq 0$.

3·7. *Non-orthogonal Covariates*

In this section we relax the assumption of orthogonality of the covariates. Let $\mathbf{X}$ denote a $n \times r$ design matrix such that the covariance matrix $var(\mathbf{X})$ has some non-zero off-diagonal entries. Let $\mathbf{U}$ be the orthogonalized matrix of covariates constructed using orthogonalization techniques such as the Gram-Schmidt algorithm. Proposition 1 states certain facts about the performance of the proposed method when the original vector of covariates $\mathbf{X}$ is replaced by $\mathbf{U}$ in the objective function $M(.)$. Specifically, it states that the proposed method still has the sparsity property as long as the penalty function is constructed based on the original $\mathbf{X}$, i.e., $\widetilde{\alpha}_y$ and $\widetilde{\alpha}_d$ are the least squares (or ridge) estimate of the parameters in the outcome and treatment models given $\mathbf{X}$, respectively. Define

$$M^u(\alpha) = \frac{1}{2n}\Big[\Big|\sum_{i=1}^n \mathbf{u}_i\widetilde{y}_i\Big| - n\alpha^\top\Big]\Big[\Big|\sum_{i=1}^n \mathbf{u}_i\widetilde{y}_i\Big| - n\alpha^\top\Big]^\top + \frac{1}{\tau}\Big[-\Big|\sum_{i=1}^n \mathbf{u}_i d_i\Big|\alpha + \sum_{i=1}^n \log(1+\exp\{\mathbf{u}_i\alpha\})\Big] \quad (7)$$

Similar to Section 3·3, let $\alpha_0^u = \alpha_{01}^u$ be the true parameter value in the model that includes $\mathbf{U}$. Let $\widehat{\alpha}^u$ be the solution to

$$\arg\min_{(\alpha^u)} M^u(\alpha^u) + \lambda\sum_{j=1}^r \frac{|\alpha_j^u|}{(\widetilde{\alpha}_{jy})^2(1+|\widetilde{\alpha}_{jd}|)^2}.$$

Because of the orthogonalization there might be mismatches between the indices of coefficients with zero value. For example, it may happen that the jth element of the vector of parameters $\alpha_j^u \neq 0$ while $\alpha_j = 0$. Therefore, if our goal is to identify the set of zero coefficients based on the covariate vector $\mathbf{X}$, the penalty function has to impose heavier penalty on coefficients that are

Table 1: Simulation results for scenarios 1 & 2. Bias, S.D. and MSE are for the treatment effect. S.D: empirical standard error. Y-fit is obtained by penalizing the outcome model via the LASSO penalty.

Method	Bias	S.D	MSE	Bias	S.D	MSE
Scenario 1.		$n = 300$			$n = 500$	
PMOE$^{\tau=0.1}$	0.031	0.633	0.402	0.036	0.431	0.187
PMOE$^{\tau=0.5}$	0.004	0.580	0.337	0.041	0.432	0.188
PMOE$^{\tau=1}$	0.007	0.598	0.357	0.040	0.431	0.188
PMOE$^{\tau=20}$	0.008	0.598	0.358	0.040	0.431	0.188
Cred. Reg.	0.004	0.780	0.608	0.024	0.585	0.342
BAC (ω=∞)	0.064	1.332	1.781	0.040	1.062	1.130
Oracle	0.005	0.593	0.351	0.030	0.430	0.186
Scenario 2.		$n = 300$			$n = 500$	
PMOE$^{\tau=0.1}$	0.027	0.556	0.310	0.008	0.419	0.175
PMOE$^{\tau=0.5}$	0.061	0.561	0.321	0.032	0.415	0.173
PMOE$^{\tau=1}$	0.123	0.587	0.360	0.065	0.424	0.184
PMOE$^{\tau=20}$	0.675	0.584	0.798	0.788	0.465	0.837
Cred. Reg.	0.010	0.765	0.582	0.027	0.595	0.355
Y-fit	0.710	0.598	0.862	0.818	0.453	0.875
BAC (ω=∞)	0.070	1.242	1.542	0.008	1.011	1.022
Oracle	0.026	0.560	0.315	0.025	0.419	0.176

significant in a model based on $\mathbf{U}$ but insignificant in the original covariate vector $\mathbf{X}$. Similarly, it may happen that the jth element of the vector of parameters $\alpha_j^u = 0$ while $\alpha_j \neq 0$. In this case, the penalty imposed on the jth covariate has to go to zero as n goes to infinity. This ensures that the penalty is not strong enough to set $\hat{\alpha}_j^u = 0$ for any finite n.

PROPOSITION 1. *Let* $\lambda_n/\sqrt{n} \to 0$ *and* $\lambda_n \to \infty$. *Then* $\Pr(\widehat{\alpha}_j^u = 0 | \alpha_j = 0) \to 1$ *as* $n \to \infty$ *and, for any finite* n, *if* $\alpha_j \neq 0$, $\widehat{\alpha}_j^u \neq 0$.

The results of this proposition relies on the choice of the penalty function. Because the penalty function is defined based on $\widetilde{\alpha}_y$ and $\widetilde{\alpha}_d$; that are the least squares (or ridge) estimate of the parameters given $\mathbf{X}$, the strength of the imposed penalty is derived by $\mathbf{X}$ and not $\mathbf{U}$. This helps us, first, to estimate $\widehat{\alpha}_j^u = 0$ when $\alpha_j = 0$ even if $\alpha_{0j}^u \neq 0$, and second, estimate $\widehat{\alpha}_j^u \neq 0$ when $\alpha_j \neq 0$ even if $\alpha_{0j}^u = 0$.

4. SIMULATION STUDIES

In this section, we study the performance of our proposed variable selection method using simulated data. We compare our results with BAC method introduced by Wang et al. (2012), the Bayesian credible region (Cred. Reg.) introduced by Wilson & Reich (2014) and outcome penalized estimator (Y-fit). Our simulation includes a scenario in which there is a weak confounder that is strongly related to the treatment but weakly to the outcome. We consider linear working models for $g()$ throughout this section.

We generate 500 data sets of sizes 300 and 500 from the following two models:

1. $D \sim \text{Bernoulli}\left(\dfrac{\exp\{0.2x_1 - 2x_2 + x_5 - x_6 + x_7 - x_8\}}{1+\exp\{0.2x_1 - 2x_2 + x_5 - x_6 + x_7 - x_8\}}\right)$

   $Y \sim \text{Normal}(d + 2x_1 + 0.5x_2 + 5x_3 + 5x_4, 4)$

2. $D \sim \text{Bernoulli}\left(\dfrac{\exp\{0.2x_1 - 2x_2 + x_5 - x_6 + x_7 - x_8}{1+\exp\{0.2x_1 - 2x_2 + x_5 - x_6 + x_7 - x_8\}}\right),$

   $Y \sim \text{Normal}(d + 2x_1 + 0.2x_2 + 5x_3 + 5x_4, 4)$

where $\mathbf{X}_k$ has a $N(1,4)$ for $k = 1, ..., 100$. Note that in the second scenario, x_2 is considered as a weak confounder. Results are summarized in Table 5; the Y-fit row refers to the estimator obtained by penalizing the outcome model using *LASSO* penalty, and the *Oracle* row refers to estimator obtained by including (x_1, x_2, x_3, x_4) in the propensity score and outcome model (6). We studies the performance of PMOE for different values of constant $\tau = 0.1, 0.5, 1$ and 20. In the first scenario there is no weak confounder and the Y-fit is omitted since the results are similar to the PMOE$^{\tau=20}$ row. The Bayesian adjustment for confounding (BAC) has been implemented using the `R` package `BEAU` with $\omega = \infty$ and the Bayesian credible region (Cred. Reg.) has been implemented using the `R` package `BayesPen` with flat prior.

The variance of the estimator in the BAC is too large due to the inclusion of spurious variables that are not related to the outcome. The PMOE and Creg. Reg. estimators, however, are unbiased and have smaller variance. In fact, PMOE has the lowest variance compared to BAC and Creg. Reg. methods regardless of the value of τ. In the second scenario, the Y-fit estimator is bias because of under selecting the confounder X_2. Also, the bias of the proposed PMOE estimator increases by increasing τ that is expected because as τ increases the proposed method should perform similarly to the outcome based variable selection methods such as Y-fit.

Table 6 presents the average number of coefficients set to zero correctly and incorrectly under the second scenario. There are four non-zero coefficients in our generative model so the number in the correct column should be 96 and in the incorrect column should be 0. This table shows that both Creg. Reg. and BAC are somewhat conservative and include some of the variables that should not be included which can be the source of the observed variance inflation in Table 5. This is mostly due to inclusion of variables that are predictor of treatment model but have no association with the outcome. Also, increasing τ in PMOE, increases the chance of setting the coefficient of X_2 to zero. The Y-fit row shows that this method is setting a nonzero coefficient to zero (i.e., coefficient of X_2) which explains the bias in Table 5. This, in fact, highlights the importance of our proposed method. Our simulation studies suggest that we should avoid large values of τ if we are concern about excluding weak confounders such as X_2.

Simulation studies presented in Appendix 2 study the performance of our covariate selection and estimation procedure when the covariates are non-orthogonal. Tables 5 & 6 show that the proposed method is still outperforming the other methods. Moreover, in Appendix 3, we study cases that the number of covariates is larger than the sample size ($r > n$). We also investigate cases where either of the working models of the propensity score or the outcome model is mis-specified. Our results show that the proposed method performs well, and outperforms Y-fit

5. APPLICATION: THE EFFECT OF LIFE EXPECTANCY ON ECONOMIC GROWTH

In this section we examine the performance of our proposed method on the cross-country economic growth data used by Doppelhofer et al. (2003). For illustration purposes, we focus on

Table 2: Simulation results for scenarios 1 & 2. Number of coefficients that are correctly or incorrectly set to zero. Y-fit is obtained by penalizing the outcome model via LASSO penalty.

Method	Correct	Incorrect	Correct	Incorrect
	$n = 300$		$n = 500$	
PMOE$^{\tau=0.1}$	95.40	0.03	95.74	0.01
PMOE$^{\tau=0.5}$	95.90	0.09	96.00	0.01
PMOE$^{\tau=1}$	96.00	0.18	96.00	0.06
PMOE$^{\tau=20}$	96.00	0.76	96.00	0.89
Cred. Reg.	92.39	0.00	92.61	0.00
BAC (ω=∞)	91	0.00	91	0.00
Y-fit	97	0.90	97	0.92

a subset of the data which includes 88 countries and 35 variables. Additional details are provided in Doppelhofer & Weeks (2009). We are interested in selecting non-ignorable variables which confound the effect of *life expectancy* (exposure variable) as a measure of population health on the *average growth rate of gross domestic product per capita in 1960-1996* (outcome).

The causal (or, at least, *unconfounded*) effect of life expectancy on economic growth is controversial. Acemoglu & Johnson (2006) find no evidence of increasing life expectancy on economic growth while Husain (2012) shows that it might have positive effect. We dichotomize the life expectancy based on the observed median, which is 50 years. Hence, the exposure variable D=1 if life expectancy is below 50 years in that country and 0 otherwise.

We select the significant covariates for the conditional mean and the treatment models using the penalized objective function (5). After covariate selection, we fit the model $\mathbb{E}[Y|s, \mathbf{x}] = \theta s + g(\mathbf{x}; \gamma)$, where θ is the treatment effect parameter (the function $g()$ assumed to be linear). Interaction or the higher order of the propensity score can be added to the outcome model if needed.

In our analysis, Y-fit refers to the case where just the outcome model is penalized using LASSO to select the significant covariates. We also implement the BAC with $\omega = \infty$ and Credible region with flat prior methods. Table 3 presents the list of variables and their estimated coefficients which are selected at least by one of the methods.

The proposed method selects 12 and 9 variables depending on the value of τ while Y-fit, BAC and Cred. Reg. select 7, 6 and 11 variables, respectively. Y-fit and PMOE$^{\tau=5}$ are under selecting some of non-ignorable confounders which barely predict the outcome. Specifically, *Population Density 1960*, and *Initial Income* are such non-ignorable confounders which are known to be significant in the economics literature. Table 4 shows that this under selection leads to a biased treatment effect estimator. BAC and Creg. Reg. are also ignoring some of the important potential confounders. Specifically, BAC does not select *Population Density 1960* and Creg. Reg. is ignoring *Initial outcome in 1960*. Doppelhofer & Weeks (2011) listed the latter two variables as potential confounders. These under selections may due the non-linear nature of the outcome or/and treatment model. Moreover, because there are strong evidence that variables such as *Oil producing Country, Land Area Near Navigable Water, and Nominal Government Share* are unlikely to confound the effect of the life expectancy on the economic growth (Doppelhofer & Weeks, 2009; Acemoglu & Johnson, 2006; Ley & Steel, 2007, 2009a,b; Magnus et al., 2010; Doppelhofer & Weeks, 2011; Eicher et al., 2011), PMOE$^{\tau=0.5}$ seems to select more reasonable covariates than PMOE$^{\tau=0.1}$.

Table 3: The economic growth data: List of significant variables. Y-fit is obtained by penalizing the outcome model via LASSO penalty.

Variable	Y-fit	BAC ω=∞	Cred. Reg.	PMOE $\tau = 0.1$	PMOE $\tau = 0.5$	PMOE $\tau = 5$
Air Distance to Big Cities	—	—	—	√	√	√
Ethnolinguistic Fractionalization	√	√	√	√	√	√
Fraction of Catholics	—	—	—	√	√	√
Population Density 1960	—	—	√	√	√	—
East Asian Dummy	√	√	√	√	√	√
Initial Income (Log GDP in 1960)	—	√	—	√	√	√
Public Education Spending Share	√	—	√	—	—	—
Nominal Government Share	—	—	—	√	—	—
Investment Price	√	√	√	—	—	—
Land Area Near Navigable Water	—	—	√	—	—	—
Fraction GDP in Mining	—	—	√	√	√	√
Fraction Muslim	—	—	√	—	—	—
Political Rights	—	√	—	√	√	√
Real Exchange Rate Distortions	√	—	—	—	—	—
Colony Dummy	—	—	√	—	—	√
European Dummy	√	—	—	—	—	—
Latin American Dummy	√	√	—	√	√	√
Landlocked Country Dummy	—	—	√	—	—	—
Oil producing Country Dummy	—	—	—	√	—	—
Land Area Near Navigable Water	—	—	√	√	—	—

To gain insight into the effect of parameter τ on the selection results, we plot the estimated coefficients α for different values of τ for given tuning parameters. In Figure 3, the blue solid (dark dashed) lines correspond to coefficients that their estimated value is (not) greater than 0.05 for the entire range of τ. Also, displays (a)–(d) correspond to tuning parameter values $\lambda = 0.0$, 0.0001, 0.001, and 0.01, respectively. Figure 3 shows that, for different values of τ, the selected set of covariates may vary slightly. For example, in Figure 3(c) where a moderate penalty function is imposed, the proposed method suggests to include *Colony* (the dashed line) to the selected set for larger values of τ. Also, as the penalty becomes stronger, the estimated coeffiecints become more stable across values of τ (Figure 3(d)). This is because, for larger values of λ, the penalty function has a more dominant rule on the variable selection than τ.

Table 4 also reports the stadard errors of the estimated effect of life expectancy using different penalization methods. The standard error of the PMOE estimator is approximated using an idea similar to Chatterjee & Lahiri (2011). Specifically, we bootstrap the sample and in each bootstrap force the components of the penalized estimator $\widehat{\alpha}$ to zero whenever they are close to zero and estimate the treatment effect using the selected covariates, i.e., we define $\widehat{\alpha}^{\dagger} = \hat{\alpha} 1(|\hat{\alpha}| > 1/\sqrt{n})$. We utilize this thresholded bootstrap method to approximate the standard error of the treatment effect. Although more investigation is required to validate the asymptotic properties of this method, we only use this standard errors to shed light on the behavior of the penalized estimators. For example, the estimators correspond to PMOE$^{\tau=5}$ and Y-fit have the lowest standard errors that may support the possibility of under-selecting important covariates. Our results sug-

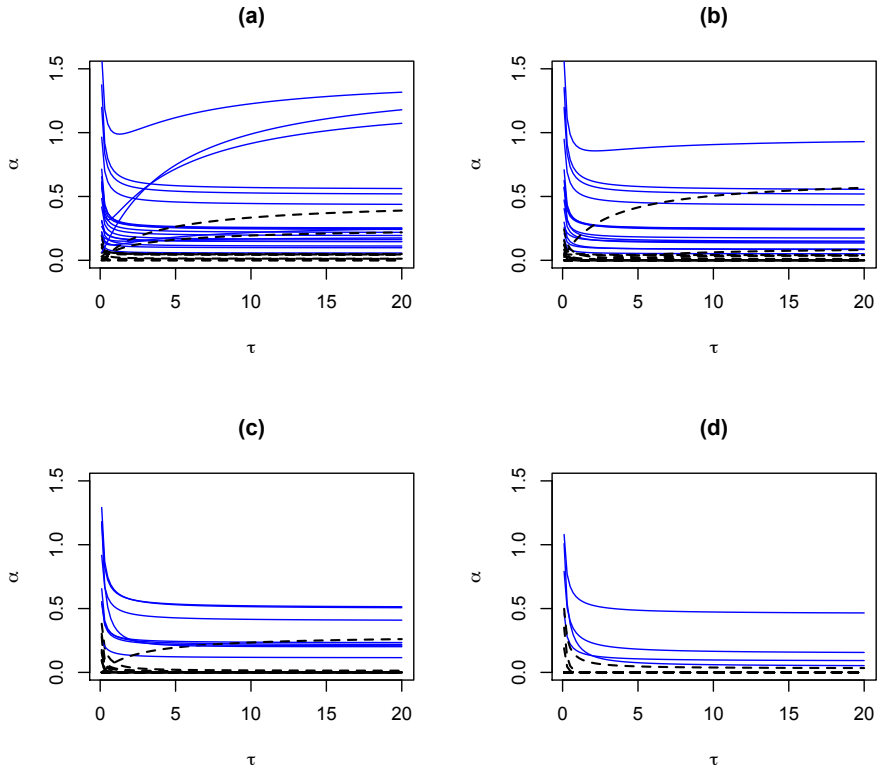

Fig. 3: Economic growth data: The plot of the estimated parameters using the proposed method given different values of parameter τ. The tuning parameter λ is fixed at 0.0, 0.0001, 0.001, and 0.01 in plots (a), (b), (c) and (d), respectively. The blue solid (dark dashed) lines correspond to coefficients that their estimated value is (not) greater than 0.05 for the entire range of τ

gests that although the effect of life expectancy is positive, it is unlikely to be significant that is consistent with Acemoglu & Johnson (2006).

6. DISCUSSION

We have established a two-step procedure for estimating an unconfounded treatment effect in high-dimensional settings. First, we deal with the sparsity by penalizing a modified objective function which considers both covariate-outcome and covariate-treatment associations. Then, the selected variables are used to form a doubly robust regression estimator of the treatment effect by incorporating the propensity score in the conditional expectation of the outcome. The selected covariates may be used in other causal techniques as well as the proposed regression method. The proposed method may also be used to identify valid instrumental variables. Specifically, one can penalize the treatment model first and record the selected variables and then apply the proposed method and identify the instrumental variables as those that are selected by the treatment model but not the proposed method.

As described in section 3·5, any covariate selection procedure which involves the outcome variable affects the subsequent inference of the selected coefficients (Lee et al., 2013; Taylor

Table 4: The economic growth data. ATE: average treatment effect; Y-fit is obtained by penalizing the outcome model via LASSO penalty.

Method	ATE	S.D.	C.I.(%95)
PMOE$^{\tau=0.1}$	0.617	0.372	(-0.127,1.361)
PMOE$^{\tau=0.5}$	0.475	0.352	(-0.229,1.179)
PMOE$^{\tau=5}$	0.454	0.340	(-0.226,1.134)
Cred. Reg.	0.820	0.386	(-0.203,1.341)
BAC (ω=∞)	0.524	0.381	(-0.228,1.286)
Y-fit	0.352	0.334	(-0.111,1.345)

et al., 2014; Tibshirani et al., 2014; Taylor & Tibshirani, 2015; Tian & Taylor, 2015; Lee et al., 2015). This is because the selected model itself is stochastic and it needs to be accounted for. Berk et al. (2012) proposes a method to produce a valid confidence interval for the coefficients of the selected model in the post-selection context. In our setting, although we do not penalize the treatment effect, the randomness of the selected model affects the inference about the causal effect parameter through confounding. Moreover, note that the oracle property of the penalized regression estimators is a pointwise asymptotic feature and does not necessarily hold for all the points in the parameter space (Leeb & Pötscher, 2005, 2008). In this manuscript, we assume that the parameter dimension (r) is fixed while the number of observation tends to infinity. One important extension to our work is to generalize the framework to cases where the tuple (n, r) tends to infinity (Negahban et al., 2009). Analyzing the convergence of the estimated vector of parameters in the more general setting requires an adaptation of restricted eigenvalue condition (Bickel et al., 2009) or restricted isometry property (Candes & Tao, 2007).

ACKNOWLEDGMENT

This research was supported in part by NIDA grant P50 DA010075 and NSF grant SES-1260782. The second and third authors acknowledge the support of Discovery Grants from the Natural Sciences and Engineering Research Council (NSERC) of Canada. The authors are grateful to Professor Dylan Small for enlightening discussion.

APPENDIX

In this Appendix, we present the assumptions and proofs of the main and other auxiliary results. We also conduct a simulation study in which either the response mean or the treatment allocation models are misspecified.

1. PROOFS OF THEOREM 1 & 2

LEMMA 1. *Suppose conditions (a) & (b) are fulfilled, further* $\lambda_n/\sqrt{n} \to 0$ *and* $\lambda_n\sqrt{n} \to \infty$. *Then there exists a local minimizer of the penalized objective function* $M_p(.)$ *for which* $\|\alpha_0 - \hat{\alpha}\| = O_p(n^{-1/2})$, *where* $\|.\|$ *is the Euclidean norm.*

Proof. Let $\alpha = \alpha_0 + b/\sqrt{n}$ for real b. Define

$$M(b) = \frac{1}{2n}\Big[\Big|\sum_{i=1}^{n}\mathbf{x}_i\widetilde{y}_i\Big| - n(\alpha_0 + b/\sqrt{n})^\top\Big]\Big[\Big|\sum_{i=1}^{n}\mathbf{x}_i\widetilde{y}_i\Big| - n(\alpha_0 + b/\sqrt{n})^\top\Big]^\top + \frac{1}{\tau}\Big[-\Big|\sum_{i=1}^{n}\mathbf{x}_i d_i\Big|(\alpha_0 + b/n) + \sum_{i=1}^{n}\log(1+\exp\{\mathbf{x}_i(\alpha_0+b/n)\})\Big] + \lambda_n\sum_{j=1}^{r}\nu_j|\alpha_{0j} + b/\sqrt{n}|$$

where $\nu_j = 1/\{(\widetilde{\alpha}_y)^2(1+|\widetilde{\alpha}_d|)^2\}$. Therefore, $\widehat{V}(b) = M(b) - M(0)$ is given by

$$\widehat{V}(b) = -\sum_{i=1}^{n}\epsilon_{i\widetilde{y}}b/\sqrt{n} + b^\top b/2 + \frac{1}{\tau}\left[-\Big|\sum_{i=1}^{n}d_i\mathbf{x}_i\Big|b/\sqrt{n} + \sum_{i=1}^{n}\log\left\{\frac{1+\exp\{\mathbf{x}_i(\alpha_0+b/\sqrt{n})\}}{1+\exp\{\mathbf{x}_i\alpha_0\}}\right\}\right] + \lambda_n\sum_{j=1}^{r}\nu_j\left[|\alpha_{0j}+b/\sqrt{n}| - |\alpha_{0j}|\right]$$

where $\sum_{i=1}^{n}\epsilon_{i\widetilde{y}} = |\sum_{i=1}^{n}\mathbf{x}_i\widetilde{y}_i| - n\alpha_0^\top$. Note that the expected value of $\epsilon_{\widetilde{y}}$ is not necessarily zero because α_0 corresponds to the modified objective function in (3). By applying the Taylor series expansion it can be written as

$$\widehat{V}(b) = -\sum_{i=1}^{n}\epsilon_{i\widetilde{y}}b/\sqrt{n} + b^\top b/2 + \frac{1}{\tau}\sum_{i=1}^{n}\left[-\epsilon_{id}b/\sqrt{n} + \frac{\exp\{\mathbf{x}_i\alpha_0\}}{[1+\exp\{\mathbf{x}_i\alpha_0\}]^2}b^\top(\mathbf{x}_i^\top\mathbf{x}_i)b/2\right] + \lambda_n\sum_{j=1}^{r}\nu_j\left[|\alpha_{0j}+b/\sqrt{n}| - |\alpha_{0j}|\right] + o_p(n)$$

where

$$\sum_{i=1}^{n}\epsilon_{id} = |\sum_{i=1}^{n}d_i\mathbf{x}_i| - \sum_{i=1}^{n}\mathbf{x}_i\frac{\exp\{\mathbf{x}_i\alpha_0\}}{1+\exp\{\mathbf{x}_i\alpha_0\}}.$$

Similar to $\epsilon_{\widetilde{y}}$, the expected value of ϵ_d is not necessarily zero. However, by (3), the expected value of $\epsilon_{\widetilde{y}} + \frac{1}{\tau}\epsilon_d$ is zero.

Using central limit theorem and laws of large numbers

$$\sum_{i=1}^{n}[\epsilon_{i\widetilde{y}} + \frac{1}{\tau}\epsilon_{id}]\frac{b}{\sqrt{n}} \xrightarrow{d} b^\top Z \qquad \frac{1}{n}\left[n\mathbb{1} + \frac{1}{\tau}\sum_{i=1}^{n}\frac{\exp\{\mathbf{x}_i\alpha_0\}}{[1+\exp\{\mathbf{x}_i\alpha_0\}]^2}\mathbf{x}_i^\top\mathbf{x}_i\right] \xrightarrow{p} \Omega(\alpha_0),$$

where $Z \sim N(0, \Sigma(\alpha_0))$ and $\mathbb{1}$ is the identity matrix. The behavior of the third element depends on the type of covariate. Let $\mathcal{A}$ be the set of indices of confounders and predictors of outcome. Thus $\mathcal{A}$ includes indices of variables we wish to select. If $j \in \mathcal{A}$, $\frac{\lambda_n}{\sqrt{n}}\widehat{\nu}_j\sqrt{n}\,[|\alpha_{0j}+b_j/\sqrt{n}| - |\alpha_{0j}|] \xrightarrow{p} 0$. Also, because $\sum_{i=1}^{n}[\epsilon_{i\widetilde{y}} + \frac{1}{\tau}\epsilon_{id}]\frac{b}{\sqrt{n}} = O_p(1)$, the order comparison of the right hand side of $\widehat{V}(b)$ implies that

$$b^\top b/2 + \frac{1}{\tau}\sum_{i=1}^{n}\left[\frac{\exp\{\mathbf{x}_i\alpha_0\}}{[1+\exp\{\mathbf{x}_i\alpha_0\}]^2}b^\top(\mathbf{x}_i^\top\mathbf{x}_i)b/2\right]$$

dominates the other terms. Therefore, for any given $\delta > 0$, there exists a sufficiently large ζ_δ such that

$$\lim_{n\to\infty}P\left\{\inf_{||b||=\zeta_\delta}\widehat{V}(b) > 0\right\} \geq 1-\delta,$$

which implies that there is a local minimizer, say $\hat{\alpha}$, that satisfies $\|\alpha_0 - \hat{\alpha}\| = O_p(n^{-1/2})$. ■

Proof of Theorem 1: **Part (a):** Lemma 1 shows that $\widehat{\alpha}_j \xrightarrow{p} \alpha_{0j}$, $\forall j \in \mathcal{A}$ which means $\Pr(j \in \mathcal{A}_n) \to 1$. The proof is complete if we show that $\forall j \notin \mathcal{A}$, $\Pr(j' \notin \mathcal{A}_n) \to 1$.

By the KKT conditions, for $\forall j' \in \mathcal{A}_n$

$$\frac{1}{\sqrt{n}}[|\sum_{i=1}^{n} x_{ij'}\widetilde{y}_i| - \sum_{i=1}^{n}\widehat{\alpha}^{\top}(x_{ij'}\mathbf{x}_i)] + \frac{1}{\tau\sqrt{n}}\left[|\sum_{i=1}^{n} x_{ij'}d_i| - \sum_{i=1}^{n} x_{ij'}\frac{\exp\{\mathbf{x}_i\widehat{\alpha}\}}{1+\exp\{\mathbf{x}_i\widehat{\alpha}\}}\right] = \frac{\lambda_n}{\sqrt{n}}\nu_{j'}.$$

Using central limit theorem and laws of large numbers

$$\sum_{i=1}^{n}[\epsilon_{i\widetilde{y}} + \frac{1}{\tau}\epsilon_{id}]\frac{b}{\sqrt{n}} \xrightarrow{d} b^{\top}Z \qquad \frac{1}{n}\left[n\mathbb{1} + \frac{1}{\tau}\sum_{i=1}^{n}\frac{\exp\{\mathbf{x}_i\alpha_0\}}{[1+\exp\{\mathbf{x}_i\alpha_0\}]^2}\mathbf{x}_i^{\top}\mathbf{x}_i\right] \xrightarrow{p} \Omega(\alpha_0),$$

where Z has a normal distribution and $\mathbb{1}$ is the identity matrix. However, for a suitable choice of λ_n (i.e., $\lambda_n/\sqrt{n} \to 0$ and $\lambda_n\sqrt{n} \to \infty$), $\frac{\lambda_n}{\sqrt{n}}\nu_{j'} \to_p \infty$. Thus, consistency of the estimator follows from Theorem 2 in Zou (2006).

Part (b): Following the results of Lemma 1 and by convexity of $\widehat{V}(b)$ and Geyer(1994), $\arg\min \widehat{V}(b) \xrightarrow{d} \arg\min V(b)$ where $\arg\min \widehat{V}(b) = \sqrt{n}(\widehat{\alpha} - \alpha_0)$. By Slutsky's theorem, $\arg\min V(b_{\mathcal{A}}) = \Omega_{11}^{-1}Z_{\mathcal{A}}$ where $Z_{\mathcal{A}} \sim N(0, \Sigma_{11})$ and $\arg\min V(b_{\mathcal{A}^c}) = 0$. This completes the proof of Normality. ■

Proof of Theorem 2: Using the triangle inequality,

$$|\zeta(\widehat{\theta}_{M_n}, M_n) - \zeta(\theta_0, M_0)| \leq |\zeta(\widehat{\theta}_{M_n}, M_n) - \zeta(\widehat{\theta}_{M_0}, M_0)| + |\zeta(\widehat{\theta}_{M_0}, M_0) - \zeta(\theta_0, M_0)|.$$

By differentiability of the $\zeta(.,.)$ function in θ, we have $\zeta(\widehat{\theta}_{M_0}, M_0) \xrightarrow{p} \zeta(\theta_0, M_0)$, so the second term on the right hand side converges to zero in probability. Also, $\forall t > 0$, we have for each n that

$$\begin{aligned}\Pr(|\zeta(\widehat{\theta}_{M_n}, M_n) - \zeta(\widehat{\theta}_{M_0}, M_0)| > t) &= \Pr(\{M_n = M_0\} \cap \{|\zeta(\widehat{\theta}_{M_n}, M_n) - \zeta(\widehat{\theta}_{M_0}, M_0)| > t\}) \\ &\quad + \Pr(\{M_n \neq M_0\} \cap \{|\zeta(\widehat{\theta}_{M_n}, M_n) - \zeta(\widehat{\theta}_{M_0}, M_0)| > t\}) \\ &\leq \Pr(M_n \neq M_0).\end{aligned}$$

We have $\Pr(M_n \neq M_0) \longrightarrow 0$ as $n \longrightarrow \infty$ by the oracle property of our procedure (Theorem 1). See also Theorem 4.2 in Wasserman and Roeder (2009). This completes the proof of weak consistency. ■

Proof of Proposition 1: Let $\alpha^u = \alpha_0^u + b/\sqrt{n}$ and $\nu_j = 1/\{(\widetilde{\alpha}_y)^2(1+|\widetilde{\alpha}_d|)^2\}$. Similar to the proof of Theorem 1, define

$$\begin{aligned}\widehat{V}^u(b) = -\sum_{i=1}^{n}\epsilon_{i\widetilde{y}}^u b/\sqrt{n} + b^{\top}b/2 + \frac{1}{\tau}\sum_{i=1}^{n}\left[-\epsilon_{id}^u b/\sqrt{n} + \frac{\exp\{\mathbf{u}_i\alpha_0^u\}}{[1+\exp\{\mathbf{u}_i\alpha_0^u\}]^2}b^{\top}b/2\right] \\ + \lambda_n\sum_{j=1}^{r}\nu_j\left[|\alpha_{0j}^u + b/\sqrt{n}| - |\alpha_{0j}^u|\right] + o_p(n)\end{aligned}$$

where $\sum_{i=1}^{n}\epsilon_{id}^u = \left[|\sum_{i=1}^{n} d_i\mathbf{u}_i| - \sum_{i=1}^{n}\mathbf{u}_i\frac{\exp\{\mathbf{u}_i\alpha_0^u\}}{1+\exp\{\mathbf{u}_i\alpha_0^u\}}\right]$ and $\epsilon_{i\widetilde{y}}^u = [|\sum_{i=1}^{n}\mathbf{u}_i\widetilde{y}_i| - n\alpha_0^{u\top}]$.

Let $\mathcal{A}$ be the set of indices of confounders and predictors of outcome based on the original covariate matrix $\mathbf{X}$. If $j \in \mathcal{A}$, $\frac{\lambda_n}{\sqrt{n}}\nu_j\sqrt{n}\left[|\alpha_{0j}^u + b_j/\sqrt{n}| - |\alpha_{0j}^u|\right] \xrightarrow{p} 0$. In other words, because $\frac{\lambda_n}{\sqrt{n}} \to 0$, the penalty function does not impose any penalty on α_j^u if $j \in \mathcal{A}$. Thus, for any finite n, $\widehat{\alpha}_j^u \neq 0\ \forall j \in \mathcal{A}$.

Now, we show that $\Pr(\widehat{\alpha}_j^u = 0; j \in \mathcal{A}^c) \to 1$ as $n \to \infty$. Let $\ell_n(\alpha_1^u, \alpha_2^u)$ be the negative log-likelihood function of (α_1^u, α_2^u) given $\mathbf{U}$. Also, we assume that, for each α_0^u, there exists function $M_1(\mathbf{u})$ such that for α in the neighborhood of α_0^u,

$$\left|\frac{\partial f(\mathbf{u}, \alpha)}{\partial \alpha}\right| \leq M_1(\mathbf{u}),$$

such that $\int M_1(\mathbf{u})d\mathbf{u} < \infty$. By the mean value theorem,

$$\ell_n(\alpha_1^u, \alpha_2^u) - \ell_n(\alpha_1^u, 0) = \left[\frac{\partial \ell_n(\alpha_1^u, \xi)}{\partial \alpha_2^u}\right]^\top \alpha_2^u,$$

for some $\|\xi\| \leq \|\alpha_2^u\|$. Also, by mean value theorem and adding and subtracting $\frac{\partial \ell_n(\alpha_1, 0)}{\partial \alpha_2^u}$, we have

$$\left\|\frac{\partial \ell_n(\alpha_1^u, \xi)}{\partial \alpha_2^u} - \frac{\partial \ell_n(\alpha_{01}^u, 0)}{\partial \alpha_2^u}\right\|^2 \leq \left[\sum_{i=1}^n M_1(x_i)\right] \|\xi\| + \left[\sum_{i=1}^n M_1(x_i)\right] \|\alpha_{01^u} - \alpha_1^u\|.$$

Let $\mathcal{B}$ be a set of indices for which $\alpha_j^u \neq 0$ while $\alpha_j = 0$. For $j \in \mathcal{B}$, $\|\xi\| \leq \|\alpha_j^u\| = O_p(1)$. Thus

$$\left\|\frac{\partial \ell_n(\alpha_1^u, \xi)}{\partial \alpha_2^u} - \frac{\partial \ell_n(\alpha_{01}^u, 0)}{\partial \alpha_2^u}\right\|^2 \leq O_p(n).$$

Because, for $j \in \mathcal{B}$, $\frac{\partial \ell_n(\alpha_{01}^u, 0)}{\partial \alpha_2^u} = O_p(n)$, we conclude that $\frac{\partial \ell_n(\alpha_1^u, \xi)}{\partial \alpha_2^u} = O_p(n)$.

Let $\ell_n^p(\alpha_1^u, \alpha_2^u) = \ell_n(\alpha_1^u, \alpha_2^u) + \lambda_n \sum_{j \in \mathcal{A} \cup \mathcal{A}^c} \widehat{\nu}_j |\alpha_j^u|$. By applying the assessed orders, we have

$$\begin{aligned}\ell_n^p(\alpha_1, \alpha_2) - \ell_n^p(\alpha_1, 0) &= \sum_{j \in \mathcal{A}^c} \{-|\alpha_j| O_p(n) + \lambda_n \widehat{\nu}_j |\alpha_j|\} \\ &= \sum_{j \in \mathcal{A}^c} \{-|\alpha_j| O_p(n) + n\lambda_n O_p(1)|\alpha_j|\}. \qquad \text{(A1)}\end{aligned}$$

Thus,

$$\ell_n^p(\alpha_{\mathcal{A}}, \alpha_{\mathcal{A}^c}) - \ell_n^p(\widehat{\alpha}_{\mathcal{A}}, 0) \pm \ell_n^p(\alpha_{\mathcal{A}}, 0) \geq \ell_n^p(\alpha_{\mathcal{A}}, \alpha_{\mathcal{A}^c}) - \ell_n^p(\alpha_{\mathcal{A}}, 0).$$

By (A1), the RHS of the above inequality is positive with probability 1 when $\lambda_n \to \infty$ as $n \to \infty$. ■

2. NON-ORTHOGONAL COVARIATES

In this section, we study the performance of our covariate selection and estimation procedure when the covariates are non-orthogonal. We generate 500 data sets of sizes 300 and 500 from the following two models:

$$D \sim \text{Bernoulli}\left(\frac{\exp\{x_1 - 2x_2 + x_5 - x_6 + x_7 - x_8}{1 + \exp\{x_1 - 2x_2 + x_5 - x_6 + x_7 - x_8\}}\right),$$

$$Y \sim \text{Normal}(d + x_1 + 0.2x_2 - x_3 + x_4, 2)$$

where $\mathbf{X}_k \sim N(1, \sqrt{2})$ for $k = 1, ..., 100$. The covariate x_2 is considered as a weak confounder.

Tables 5 & 6 show that the proposed method is still outperforming the other methods. Specifically, the variance of the estimator in the Bayesian credible region (Cred. Reg.) is large due to the inclusion of spurious variables that are not related to the outcome (Table 6). This is the cause of bias in the Cred. Reg. estimator for small sample size $n = 300$.

3. PERFORMANCE UNDER MODEL MISSPECIFICATION

In this simulation study, we want to examine the performance of our proposed method when 1) either of the outcome or treatment working models are misspecified and 2) the number of potential confounders (r) is larger then the sample size.

1. $D \sim \text{Bernoulli}\left(\frac{\exp\{0.5x_1 - x_2 + 0.5x_5 - 0.5x_6 + 0.5x_7\}}{1 + \exp\{0.5x_1 - x_2 + 0.5x_5 - 0.5x_6 + 0.5x_7\}}\right)$

Table 5: Simulation results for non-orthogonal covariates. Bias, S.D. and MSE are for the treatment effect. S.D: empirical standard error. Y-fit is obtained by penalizing the outcome model via LASSO penalty.

	$n = 300$			$n = 500$		
Method	Bias	S.D	MSE	Bias	S.D	MSE
PMOE$^{\tau=0.5}$	0.013	0.365	0.133	0.028	0.265	0.071
Cred. Reg.	0.169	0.374	0.168	0.050	0.305	0.096
Y-fit	0.265	0.375	0.211	0.190	0.359	0.169
Oracle	0.031	0.300	0.091	0.015	0.222	0.050

Table 6: Simulation results for non-orthogonal covariates. Number of coefficients that are correctly or incorrectly set to zero. Y-fit is obtained by penalizing the outcome model via LASSO penalty.

Method	Correct	Incorrect	Correct	Incorrect
	$n = 300$		$n = 500$	
PMOE$^{\tau=0.5}$	94.24	0.12	95.08	0.08
Cred. Reg.	90.03	0.00	90.04	0.00
Y-fit	97	0.71	97	0.55

$Y \sim \text{Normal}(d + x_1 + 0.2x_2 + 3x_3 + 3x_4, 2)$

2. $D \sim \text{Bernoulli}\left(\dfrac{\exp\{0.1x_1 + x_2 + 0.7\frac{x_{10}+x_9}{1+|x_8|}\}}{1+\exp\{0.1x_1 + x_2 + 0.7\frac{x_{10}+x_9}{1+|x_8|}\}}\right)$

$Y \sim \text{Normal}(d + x_1 + 0.2x_2 + 3x_3 + 3x_4, 2)$

3. $D \sim \text{Bernoulli}\left(\dfrac{\exp\{0.5x_1 - x_2 + 0.5x_5 - 0.5x_6 + 0.5x_7\}}{1+\exp\{0.5x_1 - x_2 + 0.5x_5 - 0.5x_6 + 0.5x_7\}}\right)$

$Y \sim \text{Normal}\left(d + x_1 + 2\dfrac{\exp\{0.2x_3 + 0.2x_4\}}{\exp\{0.2|x_1| + 0.2|x_2|\}}, 2\right)$

where $\mathbf{X}_k \sim N(0, 2)$ for $k = 1, ..., 550$.

In all the scenarios, we consider linear working models for the treatment and outcome. Thus, in scenarios 2 & 3, at least one of them is misspecified. Table 7 summarized the results. Y-fit refers to the estimator obtained by penalizing the outcome model using SCAD penalty.

In scenarios 1 & 2, x_2 is a non-ignorable confounder which is weakly associated with the outcome. Ignoring this variable by Y-fit method results in bias which does not go zero by increasing the sample size. Our proposed method PMOE outperforms Y-fit by increasing the chance of including all the confounders (weak or strong) in the model. Our simulation also shows that the proposed method selects important covariates even when the true generative models are non-linear. It also highlights the importance of using a double robust estimator in the estimation step to obtain a consistent treatment effect estimate when at least one of the outcome or treatment models are correctly specified.

Table 7: Performance of the proposed method when either of the outcome or treatment models are misspecified and $r > n$.

Method	Bias	S.D.	MSE	Bias	S.D.	MSE
Scenario 1.		$n = 300$			$n = 500$	
PMOE$^{\tau=0.5}$	0.010	0.304	0.092	0.005	0.222	0.049
Y-fit	0.301	0.343	0.208	0.243	0.325	0.165
Oracle	0.006	0.285	0.081	0.003	0.203	0.041
Scenario 2.		$n = 300$			$n = 500$	
PMOE$^{\tau=0.5}$	0.026	0.336	0.114	0.004	0.260	0.068
Y-fit	0.060	0.407	0.169	0.087	0.348	0.129
Oracle	0.020	0.322	0.104	0.003	0.258	0.067
Scenario 3.		$n = 300$			$n = 500$	
PMOE$^{\tau=0.5}$	0.051	0.363	0.134	0.012	0.264	0.070
Y-fit	0.091	0.373	0.147	0.112	0.311	0.109
Oracle	0.002	0.313	0.098	0.009	0.225	0.051

REFERENCES

ACEMOGLU, D. & JOHNSON, S. (2006). Disease and development: the effect of life expectancy on economic growth. Tech. rep., National Bureau of Economic Research.

ANTONIADIS, A. (1997). Wavelets in statistics: a review. *Statistical Methods and Applications* **6**, 97–130.

BANG, H. & ROBINS, J. (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics* **61**, 962–972.

BELLONI, A., CHERNOZHUKOV, V. & HANSEN, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies* **81**, 608–650.

BERK, R., BROWN, L., BUJA, A., ZHANG, K. & ZHAO, L. (2012). Valid post-selection inference. *Submitted Ann. Statist. http* .

BICKEL, P. J., RITOV, Y. & TSYBAKOV, A. B. (2009). Simultaneous analysis of lasso and Dantzig selector. *The Annals of Statistics* , 1705–1732.

BROOKHART, M. A., SCHNEEWEISS, S., ROTHMAN, K. J., GLYNN, R. J., AVORN, J. & STURMER, T. (2006a). Variable selection for propensity score models. *American Journal of Epidemiology* **163**, 1149–1156.

BROOKHART, M. A. & VAN DER LAAN, M. J. (2006b). A semiparametric model selection criterion with applications to the marginal structural model. *Computational Statistics & Data Analysis* **50**, 475–498.

CANDES, E. & TAO, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n. *The Annals of Statistics* , 2313–2351.

CHATTERJEE, A. & LAHIRI, S. N. (2011). Bootstrapping lasso estimators. *Journal of the American Statistical Association* **106**, 608–625.

CRAINICEANU, C., DOMINICI, F. & PARMIGIANI, G. (2008). Adjustment uncertainty in effect estimation. *Biometrika* **95**, 635.

DAVIDIAN, M., TSIATIS, A. & LEON, S. (2005). Semiparametric estimation of treatment effect in a pretest–posttest study with missing data. *Statistical Science* **20**, 261.

DE LUNA, X., WAERNBAUM, I. & RICHARDSON, T. (2011). Covariate selection for the nonparametric estimation of an average treatment effect. *Biometrika* **98**, 861–875.

DOPPELHOFER, G., MILLER, R. & SALA-I MARTIN, X. (2003). Determinants of long-term growth: A Bayesian averaging of classical estimates (BACE) approach. *American Economic Review* .

DOPPELHOFER, G. & WEEKS, M. (2009). Jointness of growth determinants. *Journal of Applied Econometrics* **24**, 209–244.

DOPPELHOFER, G. & WEEKS, M. (2011). Robust growth determinants. Tech. rep., CESifo working paper: Fiscal Policy, Macroeconomics and Growth.

EICHER, T. S., PAPAGEORGIOU, C. & RAFTERY, A. E. (2011). Default priors and predictive performance in bayesian model averaging, with application to growth determinants. *Journal of Applied Econometrics* **26**, 30–55.

FAN, J. & LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association* **96**, 1348–1261.

GREENLAND, S. (2008). Invited commentary: variable selection versus shrinkage in the control of multiple confounders. *American Journal of Epidemiology* **167**, 523.

HUSAIN, M. J. (2012). Alternative estimates of the effect of the increase of life expectancy on economic growth. *Economics Bulletin* **32**, 3025–3035.

KANG, J. & SCHAFER, J. (2007). Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical Science* **22**, 523–539.

LEE, J. D., SUN, D. L., SUN, Y. & TAYLOR, J. E. (2013). Exact post-selection inference with the lasso. *arXiv preprint arXiv:1311.6238* .

LEE, J. D., SUN, Y., TAYLOR, J. E. et al. (2015). On model selection consistency of regularized m-estimators. *Electronic Journal of Statistics* **9**, 608–642.

LEEB, H. & PÖTSCHER, B. (2005). Model selection and inference: Facts and fiction. *Econometric Theory* **21**, 21–59.

LEEB, H. & PÖTSCHER, B. (2008). Sparse estimators and the oracle property, or the return of Hodges' estimator. *Journal of Econometrics* **142**, 201–211.

LEY, E. & STEEL, M. F. (2007). Jointness in bayesian variable selection with applications to growth regression. *Journal of Macroeconomics* **29**, 476–493.

LEY, E. & STEEL, M. F. (2009a). Comments on jointness of growth determinants. *Journal of Applied Econometrics* **24**, 248–251.

LEY, E. & STEEL, M. F. (2009b). On the effect of prior assumptions in bayesian model averaging with applications to growth regression. *Journal of applied econometrics* **24**, 651–674.

MAGNUS, J. R., POWELL, O. & PRÜFER, P. (2010). A comparison of two model averaging techniques with an application to growth empirics. *Journal of Econometrics* **154**, 139–153.

NEGAHBAN, S., RAVIKUMAR, P. D., WAINWRIGHT, M. J., YU, B. et al. (2009). A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. In *NIPS*.

NEUGEBAUER, R. & VAN DER LAAN, M. (2005). Why prefer double robust estimators in causal inference? *Journal of Statistical Planning and Inference* **129**, 405–426.

ROBINS, J. M. (2000). Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association Section on Bayesian Statistical Science, 1999*.

ROBINS, J. M. & BRUMBACK, B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology* **11**, 550–560.

ROBINS, J. M., MARK, S. D. & NEWEY, W. K. (1992). Estimating exposure effects by modelling the expectation of exposure conditional on confounders. *Biometrics* **48**, 479–495.

ROSENBAUM, P. (2010). Causal inference in randomized experiments. *Design of Observational Studies* , 21–63.

ROSENBAUM, P. R. & RUBIN, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* **70**, 41–55.

RUBIN, D. B. (2008). For objective causal inference, design trumps analysis. *The Annals of Applied Statistics* **2**, 808–840.

SCHAFER, J. L. & KANG, J. D. Y. (2005). Discussion of "semi-parametric estimation of treatment effect in a pretest–postest study with missing data" by M. Davidian et al. *Statistical Science* **20**, 292—295.

SCHISTERMAN, E. F., COLE, S. & PLATT, R. W. (2009). Overadjustment bias and unnecessary adjustment in epidemiologic studies. *Epidemiology* **20**, 488.

SINISI, S., POLLEY, E., PETERSEN, M., RHEE, S. & VAN DER LAAN, M. (2007). Super learning: an application to the prediction of HIV-1 drug resistance. *Statistical applications in genetics and molecular biology* **6**, 7.

TAYLOR, J., LOCKHART, R., TIBSHIRANI, R. J. & TIBSHIRANI, R. (2014). Exact post-selection inference for forward stepwise and least angle regression. *arXiv preprint arXiv:1401.3889* .

TAYLOR, J. & TIBSHIRANI, R. J. (2015). Statistical learning and selective inference. *Proceedings of the National Academy of Sciences* **112**, 7629–7634.

TIAN, X. & TAYLOR, J. E. (2015). Selective inference with a randomized response. *arXiv preprint arXiv:1507.06739* .

TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B* **58**, 267–288.

TIBSHIRANI, R., TAYLOR, J., LOCKHART, R. & TIBSHIRANI, R. (2014). Exact post-selection inference for sequential regression procedures. *arXiv preprint arXiv:1401.3889* .

TSIATIS, A. A. (2006). *Semiparametric theory and missing data*. Springer Verlag.

VAN DER LAAN, M., DUDOIT, S. & VAN DER VAART, A. (2004). The cross-validated adaptive epsilon-net estimator. *UC Berkeley Division of Biostatistics Working Paper Series* , 142.

VAN DER LAAN, M., POLLEY, E. & HUBBARD, A. (2007). Super learner. *Statistical Applications in Genetics and Molecular Biology* **6**, 25.

VAN DER LAAN, M. & ROBINS, J. (2003). *Unified methods for censored longitudinal data and causality*. Springer Verlag.

VANSTEELANDT, S., BEKAERT, M. & CLAESKENS, G. (2010). On model selection and model misspecification in causal inference. *Statistical Methods in Medical Research* , 1477–0334.

WANG, C., PARMIGIANI, G. & DOMINICI, F. (2012). Bayesian effect estimation accounting for adjustment uncertainty. *Biometrics* **68**, 661–671.

WILSON, A. & REICH, B. J. (2014). Confounder selection via penalized credible regions. *Biometrics* .

ZIGLER, C. M., WATTS, K., YEH, R. W., WANG, Y., COULL, B. A. & DOMINICI, F. (2013). Model feedback in Bayesian propensity score estimation. *Biometrics* .

ZOU, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American Statistical Association* **101**, 1418–1429.

ZOU, H. & HASTIE, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B* **67**, 301–320.